

"SEEING EYE" MOUSE FOR A COMPUTER SYSTEM

Cross Reference to Related Application

This Application is a Continuation of an earlier filed co-pending Application originally entitled "SEEING EYE" MOUSE FOR A COMPUTER SYSTEM, whose title is now PROXIMITY
5 *R* DETECTOR FOR A "SEEING EYE" MOUSE, ^{ser. No.} SN 09/052, 046 and filed 30 March 1998 by Gary B. Gordon et al. ^{now U.S. Pat No. 6, 281, 882}

Reference to Related Patents

This Application is related to the subject matter described in the following two U. S. Patents: 5, 578, 813 filed 2 March 1995, issued 26 November 1996 and entitled FREEHAND IMAGE SCANNING DEVICE WHICH COMPENSATES FOR NON-LINEAR MOVEMENT; and 5, 644, 139, filed 14 August 1996, issued 1 July 1997 and entitled NAVIGATION FOR DETECTING MOVEMENT OF NAVIGATION SENSORS RELATIVE TO AN OBJECT. Both of these Patents have the same inventors: Ross R. Allen, David Beard, Mark T. Smith and Barclay J. Tullis, and both Patents are assigned to Hewlett-Packard Co. This application is also related to the subject matter described in U.S. Patent 5, 786, 804 filed 6 October 1995, entitled METHOD AND SYSTEM FOR TRACKING ATTITUDE, issued 28 July 1998, and also assigned to Hewlett-Packard Co. These three Patents describe techniques of tracking position movement. Those techniques are a component in the preferred embodiment described below. Accordingly, U.S. Patents 5, 578, 813, 5, 644, 139 and 5, 786, 804 are hereby incorporated herein by reference.

Background Of The Invention

The use of a hand operated pointing device for use with a computer and its display has become almost universal. By far the most popular of the various devices is the conventional (mechanical) mouse. A conventional mouse typically has a bottom surface carrying three or more downward projecting pads of a low friction material that raise the bottom surface a short distance above the work surface of a cooperating mouse pad. Centrally located within the bottom surface of the mouse is a hole through
25 which a portion of the underside of a rubber-surfaced steel ball (hereinafter called simply a rubber ball)

extends; in operation gravity pulls the ball downward and against the top surface of the mouse pad. The mouse pad is typically a closed cell foam rubber pad covered with a suitable fabric. The low friction pads slide easily over the fabric, but the rubber ball does not skid, but instead rolls as the mouse is moved. Interior to the mouse are rollers, or wheels, that contact the ball at its equator (the great circle parallel to the bottom surface of the mouse) and convert its rotation into electrical signals. The external housing of the mouse is shaped such that when it is covered by the user's hand it appears to have a "front-to-back" axis (along the user's forearm) and an orthogonal "left-to-right" axis. The interior wheels that contact the ball's equator are arranged so that one wheel responds only to rolling of the ball that results from a motion component of the mouse that is along the front-to-back axis, and also so that the other wheel responds only to rolling produced by a motion component along the left-to-right axis. The resulting rotations of the wheels or contact rollers produce electrical signals representing these motion components. (Say, F/B representing Forward and Backward, and L/R representing Left or Right.) These electrical signals F/B and L/R are coupled to the computer, where software responds to the signals to change by a Δx and a Δy the displayed position of a pointer (cursor) in accordance with movement of the mouse. The user moves the mouse as necessary to get the displayed pointer into a desired location or position. Once the pointer on the screen points at an object or location of interest, one of one or more buttons on the mouse is activated with the fingers of the hand holding the mouse. The activation serves as an instruction to take some action, the nature of which is defined by the software in the computer.

Unfortunately, the usual sort of mouse described above is subject to a number of shortcomings. Among these are deterioration of the mouse ball or damage to its surface, deterioration or damage to the surface of the mouse pad, and degradation of the ease of rotation for the contact rollers (say, (a) owing to the accumulation of dirt or of lint, or (b) because of wear, or (c) both (a) and (b)). All of these things can contribute to erratic or total failure of the mouse to perform as needed. These episodes can be rather frustrating for the user, whose complaint might be that while the cursor on the screen moves in all other directions, he can't get the cursor to, say, move downwards. Accordingly, industry has responded by making the mouse ball removable for easy replacement and for the cleaning of the recessed region into which it fits. Enhanced mouse ball hygiene was also a prime motivation in the introduction of mouse pads. Nevertheless, some users become extremely disgusted with their particular mouse of the moment when these remedies appear to be of no avail. Mouse and mouse pad replacement is a lively business.

The underlying reason for all this trouble is that the conventional mouse is largely mechanical in its construction and operation, and relies to a significant degree on a fairly delicate compromise about how mechanical forces are developed and transferred.

5 There have been several earlier attempts to use optical methods as replacements for mechanical ones. These have included the use of photo detectors to respond to mouse motion over specially marked mouse pads, and to respond to the motion of a specially striped mouse ball. U. S. Patent 4,799,055 describes an optical mouse that does not require any specially pre-marked surface. (Its disclosed two orthogonal one pixel wide linear arrays of photo sensors in the X and Y directions and its state-machine motion detection mechanism make it a distant early cousin to the technique of the incorporated Patents, 10 although it is our view that the shifted and correlated array [pixel pattern within an area] technique of the incorporated Patents is considerably more sophisticated and robust.) To date, and despite decades of user frustration with the mechanical mouse, none of these earlier optical techniques has been widely accepted as a satisfactory replacement for the conventional mechanical mouse. Thus, it would be desirable if there were a non-mechanical mouse that is viable from a manufacturing perspective, relatively inexpensive, reliable, and that appears to the user as essentially the operational equivalent of the conventional mouse. This need could be met by a new type of optical mouse has a familiar "feel" and is free of unexpected behaviors. It would be even better if the operation of this new optical mouse did not rely upon cooperation with a mouse pad, whether special or otherwise, but was instead able to navigate upon almost any arbitrary surface.

Summary Of The Invention

25 A solution to the problem of replacing a conventional mechanical mouse with an optical counterpart is to optically detect motion by directly imaging as an array of pixels the various particular spatial features of a work surface below the mouse, much as human vision is believed to do. In general, this work surface may be almost any flat surface; in particular, the work surface need not be a mouse pad, special or otherwise. To this end the work surface below the imaging mechanism is illuminated from the side, say, with an infrared (IR) light emitting diode (LED). A surprisingly wide variety of surfaces create a rich collection of highlights and shadows when illuminated with a suitable angle of incidence. That angle is generally low, say, on the order of five to twenty degrees, and we shall term it a "grazing" angle of

incidence. Paper, wood, formica and painted surfaces all work well; about the only surface that does not work is smooth glass (unless it is covered with fingerprints!). The reason these surfaces work is that they possess a micro texture, which in some cases may not be perceived by the unaided human senses.

IR light reflected from the micro textured surface is focused onto a suitable array (say, 16 X 16 or 24 X 24) of photo detectors. The LED may be continuously on with either a steady or variable amount of illumination servoed to maximize some aspect of performance (e.g., the dynamic range of the photo detectors in conjunction with the albedo of the work surface). Alternatively, a charge accumulation mechanism coupled to the photo detectors may be "shuttered" (by current shunting switches) and the LED pulsed on and off to control the exposure by servoing the average amount of light. Turning the LED off also saves power; an important consideration in battery operated environments. The responses of the individual photo detectors are digitized to a suitable resolution (say, six or eight bits) and stored as a frame into corresponding locations within an array of memory. Having thus given our mouse an "eye", we are going to further equip it to "see" movement by performing comparisons with successive frames.

Preferably, the size of the image projected onto the photo detectors is a slight magnification of the original features being imaged, say, by two to four times. However, if the photo detectors are small enough it may be possible and desirable to dispense with magnification. The size of the photo detectors and their spacing is such that there is much more likely to be one or several adjacent photo detectors per image feature, rather than the other way around. Thus, the pixel size represented by the individual photo detectors corresponds to a spatial region on the work surface of a size that is generally smaller than the size of a typical spatial feature on that work surface, which might be a strand of fiber in a cloth covering a mouse pad, a fiber in a piece of paper or cardboard, a microscopic variation in a painted surface, or an element of an embossed micro texture on a plastic laminate. The overall size of the array of photo detectors is preferably large enough to receive the images of several features. In this way, images of such spatial features produce translated patterns of pixel information as the mouse moves. The number of photo detectors in the array and the frame rate at which their contents are digitized and captured cooperate to influence how fast the seeing-eye mouse can be moved over the work surface and still be tracked. Tracking is accomplished by comparing a newly captured sample frame with a previously captured reference frame to ascertain the direction and amount of movement. One way that may be done is to shift the entire content of one of the frames by a distance of one pixel (corresponds to a photo

detector), successively in each of the eight directions allowed by a one pixel offset trial shift (one over, one over and one down, one down, one up, one up and one over, one over in the other direction, etc.). That adds up to eight trials, but we mustn't forget that there might not have been any motion, so a ninth trial "null shift" is also required. After each trial shift those portions of the frames that overlap each other are subtracted on a pixel by pixel basis, and the resulting differences are (preferably squared and then) summed to form a measure of similarity (correlation) within that region of overlap. Larger trial shifts are possible, of course (e.g., two over and one down), but at some point the attendant complexity ruins the advantage, and it is preferable to simply have a sufficiently high frame rate with small trial shifts. The trial shift with the least difference (greatest correlation) can be taken as an indication of the motion between the two frames. That is, it provides a raw F/B and L/R. The raw movement information may be scaled and or accumulated to provide display pointer movement information (Δx and Δy) of a convenient granularity and at a suitable rate of information exchange.

The actual algorithms described in the incorporated Patents (and used by the seeing eye mouse) are refined and sophisticated versions of those described above. For example, let us say that the photo detectors were a 16 X 16 array. We could say that we initially take a reference frame by storing the digitized values of the photo detector outputs as they appear at some time t_0 . At some later time t_1 we take a sample frame and store another set of digitized values. We wish to correlate a new collection of nine comparison frames (thought to be, null, one over, one over and one up, etc.) against a version of the reference frame representing "where we were last time". The comparison frames are temporarily shifted versions of the sample frame; note that when shifted a comparison frame will no longer overlap the reference frame exactly. One edge, or two adjacent edges will be unmatched, as it were. Pixel locations along the unmatched edges will not contribute to the corresponding correlation (i.e., for that particular shift), but all the others will. And those others are a substantial number of pixels, which gives rise to a very good signal to noise ratio. For "nearest neighbor" operation (i.e., limited to null, one over, one up/down, and the combinations thereof) the correlation produces nine "correlation values", which may be derived from a summing of squared differences for all pixel locations having spatial correspondence (i.e., a pixel location in one frame that is indeed paired with a pixel location in the other frame -- unmatched edges won't have such pairing).

6

A brief note is perhaps in order about how the shifting is done and the correlation values obtained. The shifting is accomplished by addressing offsets to memories that can output an entire row or column of an array at one time. Dedicated arithmetic circuitry is connected to the memory array that contains the reference frame being shifted and to the memory array that contains the sample frame. The formulation of the correlation value for a particular trial shift (member of the nearest or near neighbor collection) is accomplished very quickly. The best mechanical analogy is to imagine a transparent (reference) film of clear and dark patterns arranged as if it were a checker board, except that the arrangement is perhaps random. Now imagine that a second (sample) film having the same general pattern is overlaid upon the first, except that it is the negative image (dark and clear are interchanged). Now the pair is aligned and held up to the light. As the reference film is moved relative to the sample film the amount of light admitted through the combination will vary according to the degree that the images coincide. The positioning that admits the least light is the best correlation. If the negative image pattern of the reference film is a square or two displaced from the image of the sample film, the positioning admits the least light will be one that matches that displacement. We take note of which displacement admits the least light; for the seeing eye mouse we notice the positioning with the best correlation and say that the mouse moved that much. That, in effect, is what happens within an integrated circuit (IC) having photo detectors, memory and arithmetic circuits arranged to implement the image correlation and tracking technique we are describing.

It would be desirable if a given reference frame could be re-used with successive sample frames. At the same time, each new collection of nine (or twenty-five) correlation values (for collections at t_1 , t_{i+1} , etc.) that originates from a new image at the photo detectors (a next sample frame) should contain a satisfactory correlation. For a hand held mouse, several successive collections of comparison frames can usually be obtained from the (16 X 16) reference frame taken at t_0 . What allows this to be done is maintaining direction and displacement data for the most recent motion (which is equivalent to knowing velocity and time interval since the previous measurement). This allows "prediction" of how to (permanently!) shift the collection of pixels in the reference frame so that for the next sample frame a "nearest neighbor" can be expected to correlate. This shifting to accommodate prediction throws away, or removes, some of the reference frame, reducing the size of the reference frame and degrading the statistical quality of the correlations. When an edge of the shifted and reduced reference frame begins to

approach the center of what was the original reference frame it is time to take a new reference frame. This manner of operation is termed "prediction" and could also be used with comparison frames that are 5 X 5 and an extended "near neighbor" (null, two over/one up, one over/two up, one over/one up, two over, one over, ...) algorithm. The benefits of prediction are a speeding up of the tracking process by streamlining internal correlation procedure (avoiding the comparison of two arbitrarily related 16 X 16 arrays of data) and a reduction of the percentage of time devoted to acquiring reference frames.

In addition to the usual buttons that a mouse generally has, our seeing eye mouse may have another button that suspends the production of movement signals to the computer, allowing the mouse to be physically relocated on the work surface without disturbing the position on the screen of the pointer. This may be needed if the operator runs out of room to physically move the mouse further, but the screen pointer still needs to go further. This may happen, say, in a UNIX system employing a display system known as "Single Logical Screen" (SLS) where perhaps as many as four monitors are arranged to each display some subportion of the overall "screen". If these monitors were arranged as one high by four across, then the left to right distance needed for a single corresponding maximal mouse movement would be much wider than usually allowed for. The usual maneuver executed by the operator for, say, an extended rightward excursion, is to simply pick the mouse up at the right side of the work surface (a mouse pad, or perhaps simply the edge of clearing on an otherwise cluttered surface of his desk), set it down on the left and continue to move it to the right. What is needed is a way to keep the motion indicating signals from undergoing spurious behavior during this maneuver, so that the pointer on the screen behaves in an expected and non-obnoxious manner. The function of the "hold" button may be performed automatically by a proximity sensor on the underside of the mouse that determines that the mouse is not in contact with the work surface, or by noticing that all or a majority of the pixels in the image have "gone dark" (it's actually somewhat more complicated than that -- we shall say more about this idea in the next paragraph). Without a hold feature, there may be some slight skewing of the image during the removal and replacement of the mouse, owing either: (a) to a tilting of the field of view as the mouse is lifted; or (b) to some perverse mistake where frames for two disparate and widely separated spatial features imaged at very different times during the removal and replacement are nevertheless taken as representing a small distance between two frames for the same feature. A convenient place for an actual hold button is along the sides of the mouse near the bottom, where the thumb and the opposing

ring finger would grip the mouse to lift it up. A natural increase in the gripping force used to lift the mouse would also engage the hold function. A hold feature may incorporate an optional brief delay upon either the release of the hold button, detection of proper proximity or the return of reasonable digitized values. During that delay any illumination control servo loops or internal automatic gain controls would have time to stabilize and a new reference frame would be taken prior to the resumption of motion detection.

And now for this business of the pixels in the image "going dark". What happens, of course, is that the IR light from the illuminating LED no longer reaches the photo detectors in the same quantity that it did, if at all; the reflecting surface is too far away or is simply not in view. However, if the seeing eye mouse were turned over, or its underside exposed to an intensely lit environment as a result of its being lifted, then the outputs of the photo detectors might be at any level. The key is that they will be uniform, or nearly so. The main reason that they become uniform is that there is no longer a focused image; all the image features are indistinct and they are each spread out over the entire collection of photo detectors. So the photo detectors uniformly come to some average level. This is in distinct contrast with the case when there is a focused image. In the focused case the correlations between frames (recall the one over, one over and one down, etc.) exhibit a distinct phenomenon.

Assume that the spatial features being tracked mapped exactly onto the photo detectors, through the lens system, and that mouse movement were jerky by exactly the amount and in the directions needed for a feature to go from detector to detector. Now for simplicity assume also that there is only one feature, and that its image is the size of a photo detector. So, all the photo detectors but one are all at pretty much the same level, and the one detector that is not at that level is at a substantially different level, owing to the feature. Under these highly idealized conditions it is clear that the correlations will be very well behaved; eight "large" differences and one small difference (a sink hole in an otherwise fairly flat surface) in a system using nine trials for a nearest neighbor algorithm (and remembering that there may have been no motion). [Note: The astute reader will notice that the "large" difference in this rather contrived example actually corresponds to, or originates with, only one pixel, and probably does not deserve to be called "large" -- recall the earlier shifted film analogy. The only light passed by the films for this example would be for the one pixel of the feature. A more normal image having a considerably more diverse collection of pixels increases the difference to where it truly is a "large" difference.]

Now, such highly idealized conditions are not the usual case. It is more normal for the image of the tracked spatial features to be both larger and smaller than the size of the photo detectors, and for the mouse motion to be continuous, following a path that allows those images to fall onto more than one detector at once. Some of the detectors will receive only a partial image, which is to say, some detectors will perform an analog addition of both light and dark. The result is at least a "broadening" of the sink hole (in terms of the number of photo detectors associated with it) and very possibly a corresponding decrease in the depth of the sink hole. The situation may be suggested by imagining a heavy ball rolling along a taut but very stretchable membrane. The membrane has a discrete integer Cartesian coordinate system associated with it. How much does the membrane distend at any integer coordinate location as the ball rolls? First imagine that the ball is of a very small diameter but very heavy, and then imagine that the ball is of a large diameter, but still weighs the same. The analogy may not be exact, but it serves to illustrate the idea of the "sink hole" mentioned above. The general case is that the generally flat surface with sharply defined sink hole becomes a broad concavity, or bowl.

We shall term the surface produced or described by the various correlation values the "correlation surface" and will, at various times, be most interested in the shape of that surface.

We say all of this to make two points. First, the shifting shape of the concavity in the correlation surface as the seeing eye mouse moves allows interpolation to a granularity finer than the simple size/spacing of the photo detectors. We point this out, with the remark that our seeing eye mouse can do that, and leave it at that. The full details of interpolation are described in the incorporated Patents. No further discussion of interpolation is believed necessary. Second, and this is our real reason for the discussion of the preceding paragraphs, is the observation that what happens when the seeing eye mouse is picked up is that the concavity in the correlation surface goes away, to be replaced by generally equal values for the correlations (i.e., a "flat" correlation surface). It is when this happens that we may say with considerable assurance that the seeing eye mouse is air borne, and can then automatically invoke the hold feature, until after such time that a suitable concavity ("bowl") reappears.

Another method for invoking or initiating a hold feature is to simply notice that the seeing eye mouse is moving faster than a certain threshold velocity (and is thus presumably experiencing an abrupt retrace motion in a maneuver intended to translate the screen pointer further than the available physical space within which the mouse is operating). Once the velocity threshold is exceeded the motion indicating

signals that would otherwise be associated with that movement are suppressed until such time as the velocity drops below a suitable level.

Brief Description Of The Drawings

Figure 1 is a simplified pictographic cut-away side view of a prior art imaging and navigation arrangement;

Figure 2 is a bottom view of a mouse constructed in accordance with the invention;

Figure 3 is a side perspective view of a mouse constructed in accordance with one aspect of the invention; and

Figure 4 is a simplified side cut-away view of a proximity sensor in the base of the mouse of Figures 2 and 3 and used to automatically activate a hold feature;

Figure 5 is a simplified flow chart describing an aspect of internal seeing eye mouse operation related to the operation of the hold feature when used in conjunction with a feature called prediction;

Figure 6 is a simplified portion of a modification of the flow chart of Figure 5 and illustrates the velocity detection method of invoking the hold feature; and

Figure 7 is a perspective view of a plotted correlation surface that has good concavity.

Description Of A Preferred Embodiment

Refer now to Figure 1, wherein is shown a simplified representation of a cut-away side view of a prior art imaging and navigation arrangement 1 that is generally of the type described by the incorporated Patents. An LED 2, which may be an IR LED, emits light which is projected by lens 3 (which instead of being separate may be an integral part of the LED's package), through orifice 13 in bottom surface 6 and onto a region 4 that is part of a work surface 5. The average angle of incidence is preferably within the range of five to twenty degrees. Although it has been omitted for clarity, the orifice 13 might include a window that is transparent for the light from LED 2, and which would serve to keep dust, dirt or other contamination out of the innards of the seeing eye mouse. Work surface 5 might belong to a special object, such as a mouse pad, or more generally, it will not, and might be the surface of nearly anything except smooth glass. Examples of suitable materials include, but are not limited to, paper, cloth, laminated plastic tops, painted surfaces, frosted glass (smooth side down, thank you), desk pads, real

wood, fake wood, etc. Generally, any micro textured surface having features whose size falls within the range of 5 to 500 microns will do.

5 The illumination of micro textured surfaces is most effective when done from the side, as this accentuates the pattern of highlights and shadows produced by surface height irregularities. Suitable angles of incidence for illumination cover the range of about five to twenty degrees. A very smooth or flat surface (e.g., one that has been ground and polished) having simple variations in reflectivity owing to (micro scale) compositional variation works, too. In such a case (and assuming that it can be guaranteed) the angle of incidence for the illumination could approach ninety degrees, since the urge to create shadows goes away. However, such a smooth yet micro detailed surface is not what we would
10 ordinarily think of when we say "arbitrary surface", and a seeing eye mouse intended for use on an "arbitrary surface" that is more likely micro textured would work best if equipped to provide a grazing angle of incident illumination.

20 An image of the illuminated region 4 is projected through an optical window 9 in package portion 8a of an integrated circuit and onto an array 10 of photo detectors. This is done with the aid of lens 7. The package portion 8a might also dispense with separate window 9 and lens 7 by combining them into one and the same element. The photo detectors may comprise a square array of, say, 12 to 24 detectors on a side, each detector being a photo transistor whose photo sensitive region is 45 by 45 microns and of 60 microns center to center spacing. The photo transistors charge capacitors whose voltages are subsequently digitized and stored in a memory. The array 10 is fabricated onto a portion of an integrated circuit die 12 affixed by an adhesive 11 onto package portion 8b. What is not shown are any of the details of how the integrated circuit is held in place (probably by a printed circuit board), the shape or composition of the lenses, or of how the lenses are mounted; it is clear that those things are doable in a conventional manner. It is also clear that the general level of illumination of region 4 may be controlled by noticing the output levels of the photo detectors and adjusting the intensity of light issuing from the
25 LED 2. This could be either continuous control or pulse width modulation, or some combination of both.

Once again, the reader is reminded that the details of the motion sensing operation are thoroughly described in the incorporated patents (and briefly described in the Summary); accordingly, they need not be repeated here.

Refer now to Figure 2, which is a bottom view of a mouse 14 constructed in accordance with the invention. In short, this bottom view of this particular seeing eye mouse 14 looks very similar to the bottom view of a particular conventional mouse from Hewlett-Packard Co., to wit: the C1413A. The major difference is that where there ought to be a ball there is a protective lens or window 16 that is transparent to IR light. This is the omitted transparent window in orifice 13 that was mentioned in the description of Figure 1. Also missing is the usual rotatable annulus that serves as a removable retainer to allow access to the ball for cleaning or replacement. What is shown in the figure is the underside 15 of the mouse 14 (corresponds to 6 in Figure 1), low friction glides 19 and connecting cable 17 with its strain relief 18. Of course, our seeing eye mouse 14 could be a cordless mouse, as well, with an optical or radio communication link to the computer.

Refer now to Figure 3, wherein is shown a side perspective view of a mouse 14 constructed in accordance with one aspect of the invention. That aspect of the invention is the hold feature. The hold feature is an aspect of seeing eye mouse operation that suspends the production of movement information or signals to the computer when it is determined that the mouse is not suitably proximate to the work surface whose spatial features are being tracked. This allows the seeing eye mouse to be picked up, translated and set back down, or, as we shall term such an operation, "swiped" across the work surface.

In particular, the seeing eye mouse 14 in Figure 3 includes at least one hold button 24 located in side skirt 20 near the bottom surface 15 so as to be beneath the right thumb or the left ring finger, depending upon which hand the operator is using. There may be another symmetrically located button on the other side (not shown) that would contact either the left thumb or the right ring finger.

The mouse 14 conventionally includes a surface 21 which nestles in the palm of the hand, and first and second "regular" mouse buttons 22 and 23 that are actuated by the index and middle fingers. These operate in their normal fashion.

Button or buttons 24 are activated by a natural increase in the gripping force needed to pick the mouse 14 up during a swipe. When one or both of these button are pressed the hold feature is activated. For the duration of the hold the sending of motion signals to the computer is suspended. When the hold is over (the buttons are released) a new reference frame is taken before any new motion signals are sent to the computer. This allows swiping, and has the advantage that the user has the ability to expressly force the onset of the hold feature.

The hold feature could also be automatically activated by the action of a separate proximity sensor on the bottom of the mouse. This is what is shown in Figure 4, where a shouldered aperture 26 in the base 6 receives a shouldered plunger 25 made captive by the lever arm of a switch 28 above. The switch 28 is activated by movement of the plunger 25, such that when the plunger moves significantly in the direction of arrow 27 the hold feature is activated. The exact nature of the separate proximity sensor is a matter of choice, and while it could be as simple as the micro switch 28 operated by the weight of the mouse through the plunger 25, other, non-mechanical, methods are also possible.

Yet another way to automatically activate and deactivate the hold feature is to examine the nature of the digitized data of the array 10 of photo detectors. When the outputs of the photo detectors become sufficiently uniform it may be surmised that there is no longer an image with variations projected onto the array 10 of photo detectors. This uniformity will reveal itself by producing a correlation surface that is flat, or nearly so. Rather than separately detecting uniform levels (which would use hardware not otherwise present), we prefer instead to examine the shape of the correlation surface, (which surface we need for other reasons, anyway). The most probable cause of a flat correlation surface is that the mouse has been picked up. This mode of operation may require that there be a fairly narrow depth of field, lest there occur undue delay in activating the hold. Such delay could produce minor artifacts in screen pointer movement. These might include slight unintended screen pointer movements owing to tilting of the mouse as it is either picked up or replaced. As long as activating the hold feature (however done, whether manually or automatically) forces acquisition of a new reference frame before resuming the production of motion signals, there should be no danger of producing a spurious indication resulting from the combination of old data with some new data that just accidentally looks like a proper small motion in some inappropriate direction. However, with mere uniform level detection (of, say, a sample frame) it may be difficult to guarantee that while in motion in the air there occur no optical effects (a reflection of a bright source) that would confuse the algorithm. It will be appreciated that the shape of the correlation surface is a much more reliable indicator. All of that said, it must still be remembered that the steering, as it were, of the screen pointer is an incrementally driven servo-like operation performed by a human being; if the screen pointer isn't there yet, just keep moving the mouse as needed! Small perturbations during swiping are not fatal, and may not even be particularly noticeable, depending upon the specific application being performed.

Refer now to Figure 5, wherein is shown a flow chart 29 that describes an aspect of seeing eye mouse operation involving the hold and prediction properties. We may assume that there is some start condition or location 30, from which is reached step 31: ACQUIRE A REFERENCE FRAME. This refers to illuminating the LED 2 and storing a collection of digitized photo detector values into an array of memory (not shown). The next step 32 is ACQUIRE A SAMPLE FRAME. This refers to the same actions, except that the data is stored in a different array of memory, and may reflect mouse motion relative to where it was when step 31 was performed. At step 33, COMPUTE CORRELATION VALUES, the nine (or perhaps twenty-five) correlation values are quickly computed by some heavy duty dedicated arithmetic hardware assisted by automatic address translation and a very wide path out of the memory arrays. At step 34, IS THE CORRELATION SURFACE SUITABLY CONCAVE?, the nature of the correlation surface described by the collection of correlation values computed in step 33 is examined. We want to know if it is shaped like a bowl, and if so, "how much water will it hold," so to speak.

If the shape of the correlation surface is a good bowl, then path 36 takes us to the optional step 37: IS THE HOLD BUTTON PRESSED?; more about that in the next paragraph. Otherwise, we have a flat correlation surface, or a "bad bowl," and proceed along path 35 to optional step 42, DELAY. There are several possible causes for this exit from qualifier 34: e.g., extreme velocity, a suddenly featureless work surface, and, an airborne mouse. In the absence of an explicit HOLD button, we will rely upon exit path 35 to provide proper seeing eye mouse behavior by suppressing motion signals to the computer during the airborne portion of a swiping operation.

If the seeing eye mouse does have a HOLD button, then optional qualifier 37 is present, and it is there that the status (pressed or not) of the HOLD 24 button is determined. The case where it is pressed is treated the same as that for a bad bowl at qualifier 34. That is, path 38 is taken, which also leads to optional step 42.

Optional step 42 provides a delay which may be useful in several ways. First, if there is a swipe in progress, then it takes some time, and by not imagining during that time some battery power can be saved. Also, suppose that the nature of the delay is slightly more complex than a pause in the motion of a moving finger on the flow chart. Suppose that the ACQUIRE REFERENCE FRAME step 31 were influenced by there having been a delay at step 42, in that part way through the delay an illumination level

control operation is initiated. This could allow time for re-adjustment of illumination levels, and so forth. Whether or not there is a DELAY at optional step 42, path 43 leads back to step 31, where another motion detection cycle begins.

5 To resume, path 39 leads to step 40: PREDICT SHIFT IN REFERENCE FRAME. As mentioned above, it is generally not necessary to obtain and maintain actual velocities in X and Y, and time interval information, to find the displacement needed for prediction. One can imagine measurement environments where that might be needed, but the one shown here is not one of them. Instead, the predicted shift can be taken as the amount of movement corresponding to the correlation at the preceding step 34.

10 The next step 44 is OUTPUT ΔX & ΔY . It is here that we take note of how much mouse motion there has been since the last measurement cycle. The amount of shift needed to attain correlation is the desired amount. These values may be found by noticing which comparison frame actually correlated (assuming no interpolation). These "raw" ΔX and ΔY motion values may be accumulated into running values that are sent to the computer at a lower rate than that at which the raw values of step 44 are produced.

15 At qualifier 45 we ask if we NEED A NEW REFERENCE FRAME?. If the answer is YES, then path 46 leads to step 48: STORE PRESENT SAMPLE FRAME IN REFERENCE FRAME. (A little thought will confirm that this re-use of the sample frame cooperates with not having to maintain actual velocities and time intervals for the prediction process. If we took a separate new reference frame it would complicate a lot of things, and would probably force the use of $D=RT$ -- i.e., the distance formula -- for prediction.)

20 We need a new reference frame when there has been enough shifting of it, owing to predictions, that not enough of it overlaps the comparison frames for reliable correlations. Somewhere in the range of three to five shifts (that do not retrace themselves) is about the limit for a 16 X 16 reference frame.

25 If the answer to qualifier 45 is NO, and we do not need to replace the reference frame, then path 47 takes us to the same step 49 as does the path leading from step 48. Step 49, SHIFT REFERENCE FRAME, performs the actual permanent shift of the values in the memory array representing the reference frame. The shift is by the prediction amount, and data shifted away is lost. Following the shifting of the reference frame path 50 returns to step 32, ACQUIRE A SAMPLE FRAME, where the next measurement cycle begins.

Refer now to Figure 6, wherein is shown a simplified flow chart segment 50 that shows how to replace step 44 of the flow chart 29 in Figure 5 with steps 51-55. The purpose for doing this is similar to the various manners of hold operation already described, and may be used in conjunction therewith, or instead thereof. The general idea of the modification represented by Figure 6 is to outfox the computer by either not sending any updated information by skipping step 55A or (optionally, with step 55B) sending zeros for ΔX and ΔY , even when that is not true. This is done whenever step 52 ascertains that the rate of mouse motion exceeds, say, three to six inches per second. For a given seeing eye mouse such a limit is easily expressed as a displacement by a certain number of pixels within some number of measurement cycles, assuming that the measurement cycle rate is fast compared to normal mouse motion. The idea is that normal casual mouse motion probably will not require either a new nearest neighbor reference frame (let alone a maximally shifted one for 5X5 near neighbor operation) every measurement cycle for some large (say, ten to twenty-five) number of consecutive measurement cycles. For if that were the case, the seeing eye mouse would be operating on the hairy edge of the hold mode via a NO answer to qualifier 34 and path 35. (According to the assumption, any higher velocity will result in loss of correlation!) That is, the expectation is that taking a new reference frame is normally much less frequent. Of course, it could happen that the velocity of the mouse is really high, and path 35 gets used, anyway. That is as it should be. But if the measurement cycle rate is not sufficiently high with respect to normal expected mouse motion, then it might not be appropriate to use the technique of Figure 6.

Step 51 represents anything in the old step 44 over and above the actual communication to the computer of the values ΔX and ΔY . A tricky example of this difference might be an internal accumulation of motion that has not yet be dispatched to the computer, owing to a higher internal motion measurement cycle rate for the seeing eye mouse than the information exchange rate with the computer. Now, it may well be the case that in some systems this accumulated information is used for internal mouse purposes other than strictly for keeping the computer informed. If so, then it would need to be preserved, for all that qualifier 52, path 53 (and bypassed step 55A) need to accomplish is NOT tell the computer there has been motion; we want to fool the computer but without making the mouse lose its mind.

It will be noticed that if such an accumulation were allowed to continued during a rapid retrace, intended to mimic picking the mouse up, the computer might still win in the end when the velocity drops to normal amounts and the accumulation is finally sent; the screen cursor could snap to the correct

location, anyway, depending upon how the overall system works. In such a case a separate set of accumulations should be maintained, with those for the computer remaining in bypassed step 55A.

Of course, it may be the case that there is no internal use by the mouse of accumulated ΔX and ΔY , other than to send it to the computer. In that case nothing needs to be done, other than to leave that accumulation in bypassed step 55A. It is also possible that in the mouse there simply are no accumulations to cause such concerns; say, any such accumulations were done by software in the computer.

Finally, refer now to Figure 7. It is a plot of a near neighbor (5 X 5) correlation surface having a suitable concavity. The two horizontal axes represent the X and Y axes of mouse motion; the units indicated along the axes are pixels. Drawn onto the plane of the axes are smoothed and interpolated contour lines intended to further indicate the shape of the correlation surface directly above. The vertical axis is a measure of correlation expressed in essentially arbitrary units.